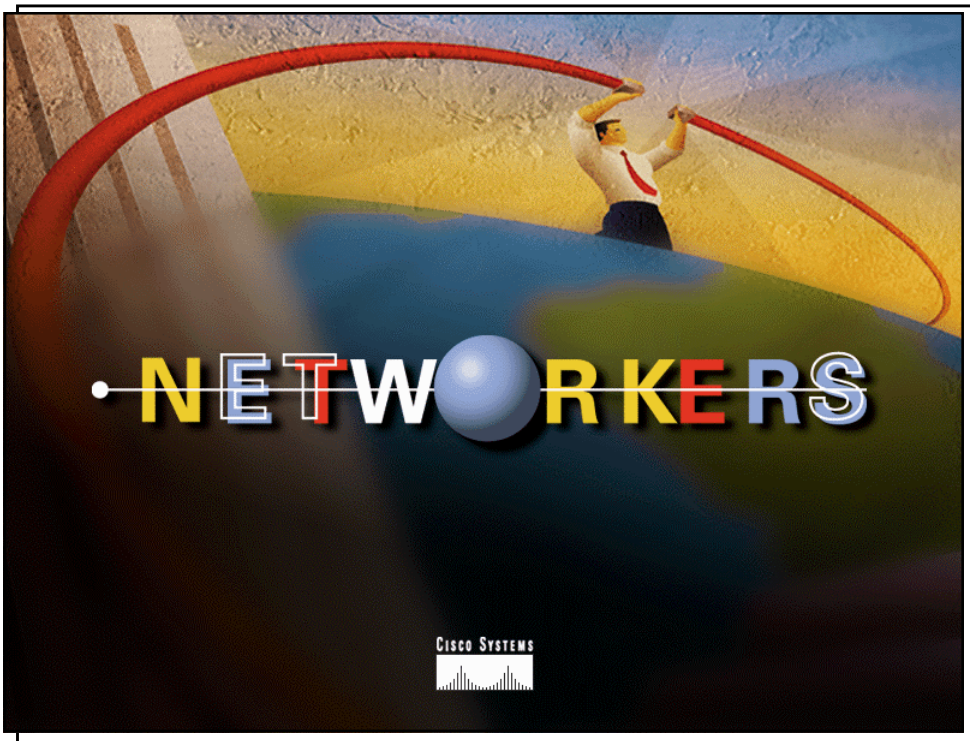

Introduction

- Traffic Management
- Applications and Transports
- So What Are the Issues for:
  - TCP
  - Voice on IP
  - Video (Broadcast and Teleconferencing)

Let's Talk About Traffic Management

- Why it is a concern
- What the guiding principles are
- What tools are available
- What can be accomplished using those tools
- What cannot be accomplished

Why Traffic Management Is a Concern

- Needs of certain applications
  - Mail? Web? Transaction processing?
- Opportunities with certain transports

www.cisco.com

Guiding Principles for Traffic Management

- We want to achieve
  - Predictability
  - Reliability
  - Availability
- In a network that
  - Keeps intelligence at the edges
  - Scales to necessary sizes and bandwidths
  - Minimizes complexity
  - Uses cost-effective technologies

www.cisco.com

What Tools Are Available for Traffic Management

- Traffic path control
- Queue depth management
- Queue rate management
- Permission to use a link

www.cisco.com

How Well Will Traffic Management Do?

- We know we can do this:
  - Management of latency
  - Management of bandwidth
- What cannot be accomplished
  - Creation of bandwidth that otherwise would not be there

www.cisco.com

Primarily a WAN IP Talk

- IP is the dominant Internet protocol
- TCP is the dominant data transport
  - 95% of Internet traffic uses TCP
- Voice is a growing market
  - But beware of hype
- Heterogeneous link layers

WAN Protocol Breakdown

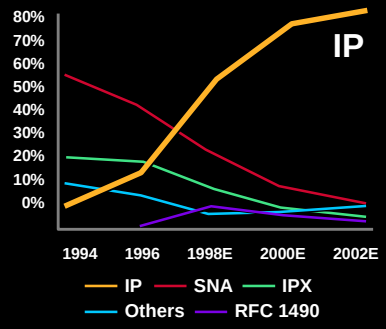

Source: Gartner Group Study, March 1997

www.cisco.com

This Is What You Need to Understand:

- TCP-based applications, voice, and video can be managed well with a little planning

www.cisco.com

Parekh and Gallagher's Paper

- INFOCOMM '93
- One must have at most a predictable amount of traffic in the network
- One must have predictable traffic delay in each network element
- Given these, end-to-end delay of a host-to-host message is predictable

www.cisco.com

Definition of “Predictable”

- Does not mean
  - “Fixed”, “Invariant”, or “Zero”
- Means that it has a
  - Mean value
  - Statistical distribution
  - Upper bound

Predictable Amount of Traffic in the Network

- The source must pace traffic initiation so that standing queues are bounded
  - Queues form when arrival rate exceeds departure rate
- When congestion (too many messages in one queue) sets in:
  - Sources must not increase their rate
  - Ideally, sources decrease their rate

Examples of Source Predictability

- TCP will keep at most a certain amount of traffic in flight
  - We say it is “elastic”—rate is proportional to latency
- Voice will send only and exactly as fast as the coding algorithm permits
  - We say it is “inelastic”

Predictable Packet Treatment in Routers and Switches

- Transit latency must be within limits acceptable to the application
- Variation in transit latency must be within limits acceptable to the application
- No stream may be locked out apart from administrative policy
- Applicable policy must be observed

Examples of Unpredictability

- Queues change rapidly enough that the distribution cannot be described
- Discards happen frequently enough that there is effectively no upper bound on delivery time

Examples of Predictability

- Classes of queues get sufficient service that ultimate arrival is timely and normal

  "Timely" is an application concept…

Quality of Service Issues in Traffic Management

- Predominantly TCP traffic
- Some specific applications
- Voice/video traffic

www.cisco.com

Backbone Traffic Mix

Transport Breakout

TCP Applications

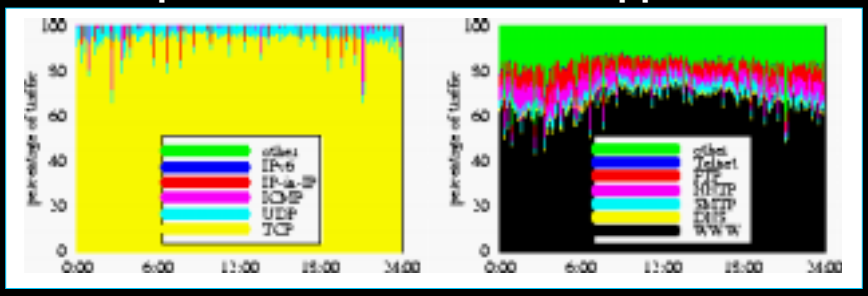

Source: MCI/NSF OC-3MON via http://www.nlanr.net, 1998

www.cisco.com

TCP Technology Issues

- Single drops communicate from network to sending host
  - "You need to slow down"
- Multiple drops in round trip trigger time-outs
  - "Something bad happened out here"

www.cisco.com

Behavior of a TCP Sender

- Sends as much as credit allows
- Starts credit small
  - Avoid overloading network queues
- Increases credit exponentially
  - To gauge network capability

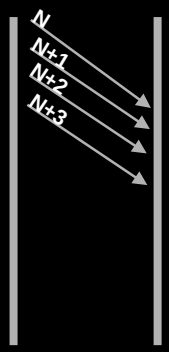

www.cisco.com

Behavior of a TCP Receiver

- When in receipt of "next message," schedules an ACK
- When in receipt of something else, acknowledges all it can immediately

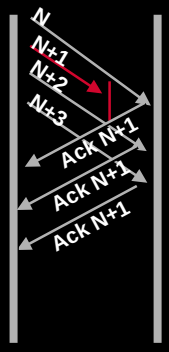

www.cisco.com

Sender Response to ACK

- If ACK acknowledges something
  - Update credit and send
- If not, presume it indicates a lost packet
  - Send first unacknowledged message right away
  - Halve current credit
  - Increase linearly to gauge network throughput

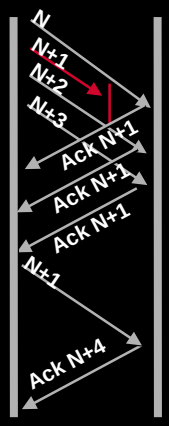

www.cisco.com

Multiple Drops in TCP

- In the event of multiple drops within the same session:
  - Current TCPs wait for time-out
  - Selective acknowledge may work around (but see INFOCOM '98)
  - New Reno "fast retransmit phase" takes several RTTs to recover

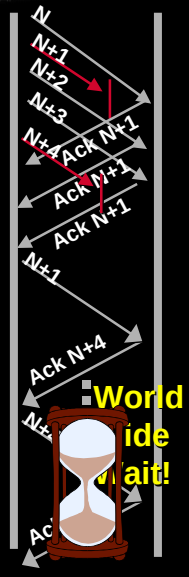

www.cisco.com

Remember Parekh and Gallagher

- One must have at most a predictable amount of traffic in the network
- One must have predictable traffic delay in each network element
- Given these, end-to-end delay of a host to host message is predictable

www.cisco.com

How Can We Make TCP in a Network Act Predictably?

- Predictable amount of traffic in the network:
  - Well-written TCP implementations manage their rates to the available bandwidth
- Router needs to
  - Provide predictable treatment of packets
  - Queue delay and drop characteristics

www.cisco.com

Fundamental FIFO Queue Management Technologies

- Tail drop
  - Network standard behavior
  - Causes session synchronization when waves of traffic experience correlated drops
- Random Early Detection (RED)
  - Random drops used to desynchronize TCP sessions and control rates

www.cisco.com

Session Synchronization

- Session synchronization results from synchronized losses
- Tail drop from waves of traffic synchronizes losses

www.cisco.com

Effect of Random Early Detection

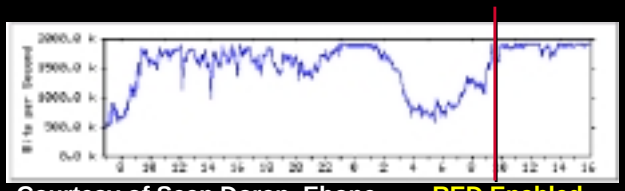

Courtesy of Sean Doran, Ebone RED Enabled

- One day, below 100% throughput
  - Simple FIFO with tail drop
- Starting 10:00 second day, 100% throughput
  - Random Early Detection enabled

www.cisco.com

Was That a Fluke?

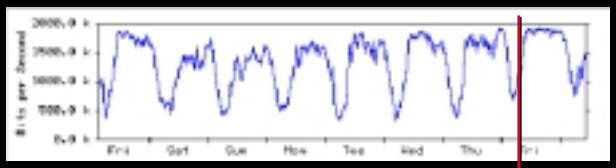

Courtesy of Sean Doran, Ebone RED Enabled

- No, here's what happened that week…
- Session synchronization reduced throughput until RED enabled

www.cisco.com

FIFO Traffic Timings

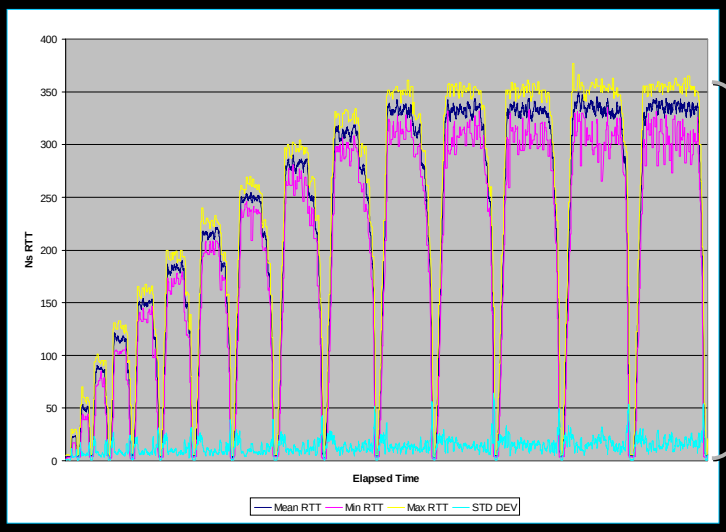

Mean Latency Correlates with Maximum Queue Depth

www.cisco.com

RED Traffic Timings

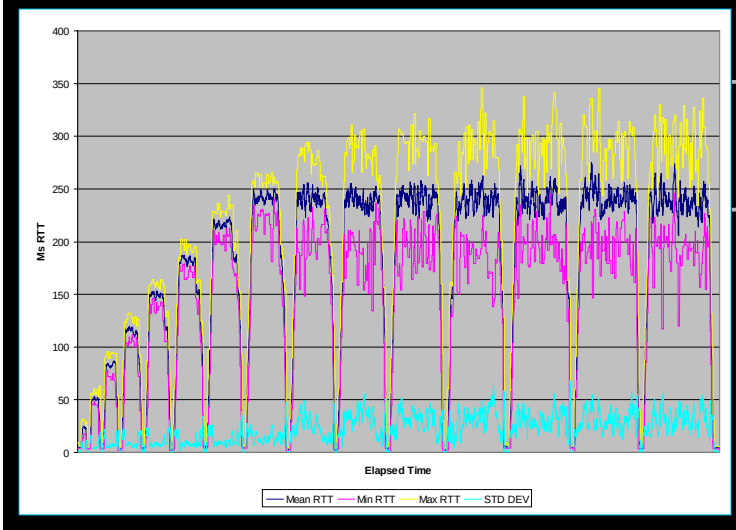

Additional Capacity to Absorb Bursts

Mean Latency Correlates with Minimum Drop Threshold

www.cisco.com

Therefore—TCP QoS Definition:

- Normally at most one drop per round trip
- Mean variation in latency bounded by predictable network

TCP Flow Statistics

- >90% of sessions have ten packets each way or less
  - Transaction mode (mail, small web page)
- >80% of all TCP traffic results from <10% of the sessions, in high rate bursts
  - It is these that we worry about managing

An Interesting Common Fallacy About RED

- "RED means you will have more drops"
  - Statement derives from observed statistics
- RED means that you will have
  - Closer to 100% utilization of your line
  - Less average delay per packet
- But queuing theory?
  - As a line approaches 100% utilization, drops will increase, even though served load increases

www.cisco.com

TCP Traffic Management Issues

- Applications
  - Often have site-specific policy associated with them
  - Traffic often identifiable by port numbers
- Sites
  - Generally identifiable by address prefix or interface traffic is received on

www.cisco.com

TCP Bandwidth Policy Questions to Answer

- Particular site or application wants at **least** a certain bandwidth
- Particular site or application wants at **most** a certain bandwidth
- Particular site or application wants to **average** about a certain bandwidth

www.cisco.com

This Is Where "Classes" Come In

- Classes can be for:
  - Voice
  - Important application/site
  - Unimportant application/site
  - Assuring at least a rate
  - Limiting to a rate

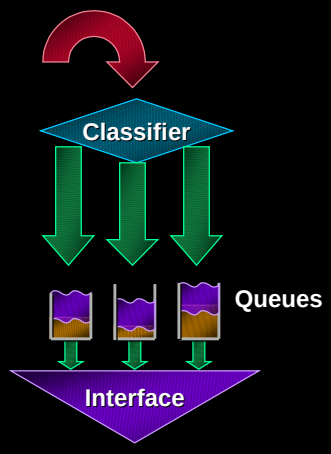

www.cisco.com

Some Class of Traffic Wants at Least a Certain Bandwidth

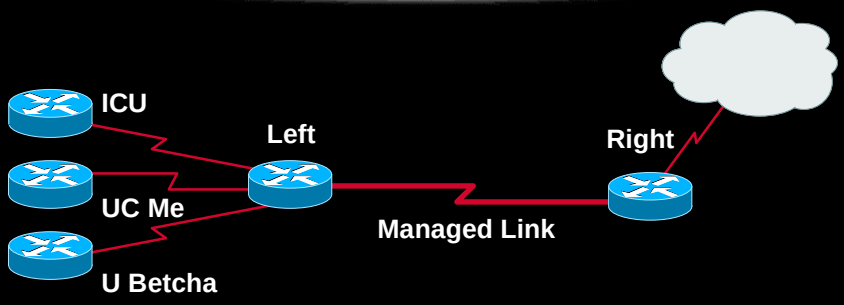

- Example:
  - Several organizations share cost of link
  - Distribute bandwidth proportional to fiscal responsibility

www.cisco.com

Class of Traffic Wants at Most a Certain Bandwidth

- Traffic shaping
- Similar queuing technology to class-based weighted fair queuing
- Rate assigned to
  - Interface or sub-interface
  - Frame Relay circuit
  - ATM virtual channel (in hardware)

www.cisco.com

Examples of Rate Control

- **Intranet exposure**

  Limit rate of web surfing outside the company

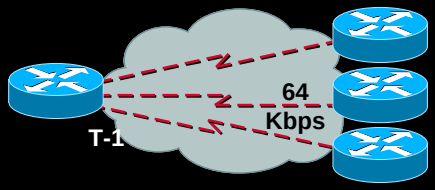

- **Frame Relay network**

  Access rate exceeds PVC rate—limit rate to rate of PVC

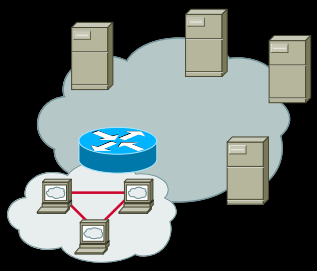

www.cisco.com

Some Class of Traffic Wants to Average a Certain Bandwidth

- **Service provider or large enterprise model**
- **Designed for**

  Cost containment

  Managed response to conflicting demands

www.cisco.com

Marking TCP Traffic at Edge

- A useful technique:
- Mark traffic at a network edge with simple classifier
- This allows network to
  - Do the right thing without having to fully classify everywhere
  - Use more effective markings

www.cisco.com

Serving TCP Traffic with the Assured Service

- Presumes service level agreement
  - Flat rate for traffic meeting a rate/burst profile
  - Usage charging for traffic out of profile
- Drop management (weighted RED)
  - All traffic subject to loss
  - Traffic out of profile much more subject to loss
  - Enhances ISP traffic engineering
    - (Good for service provider and consumer)

www.cisco.com

Structure of Presumed Service Level Agreement

- Up to rate over interval is “in profile”
- Traffic within profile gets some guarantees
- Traffic out of profile has no guarantees
  - Potentially dropped by WRED at bottleneck
  - Usage pricing of overage

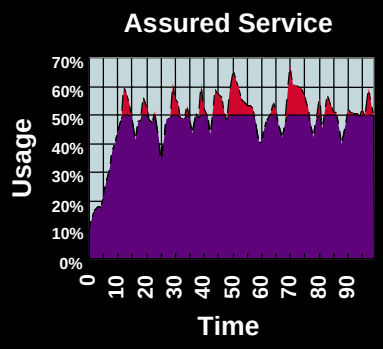

www.cisco.com

Best Effort Service in Simple IP Networks

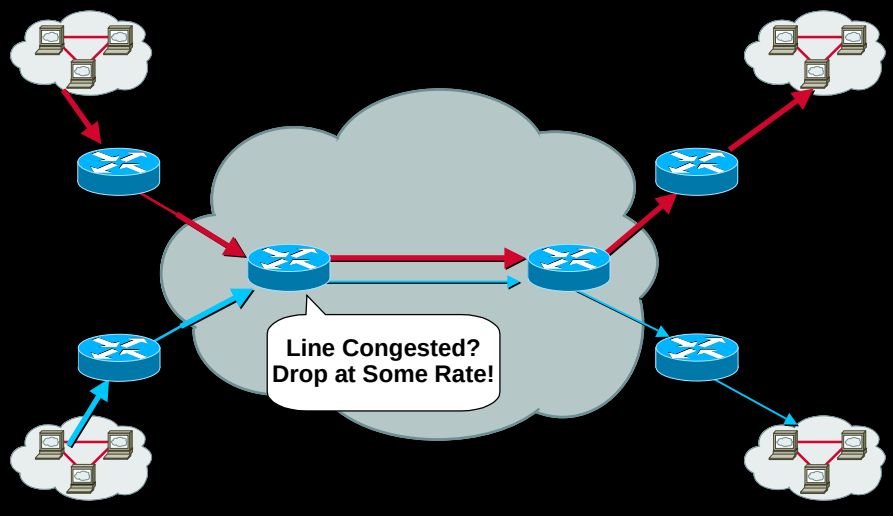

www.cisco.com

Assured Service in Simple IP Networks

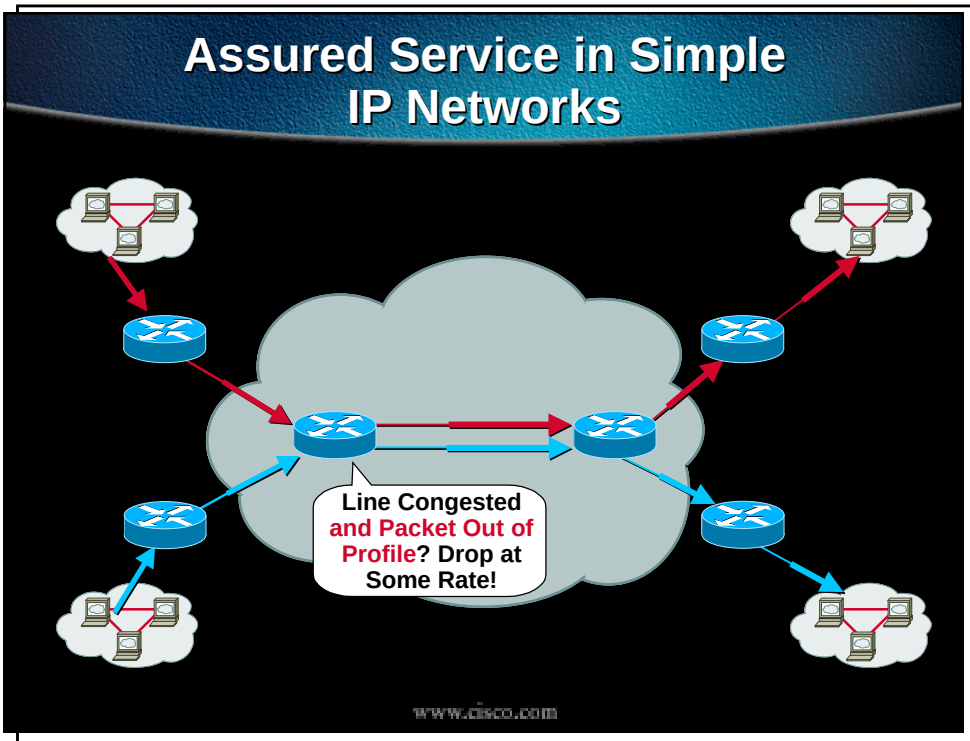

Best Effort Service in an ATM-Based Network

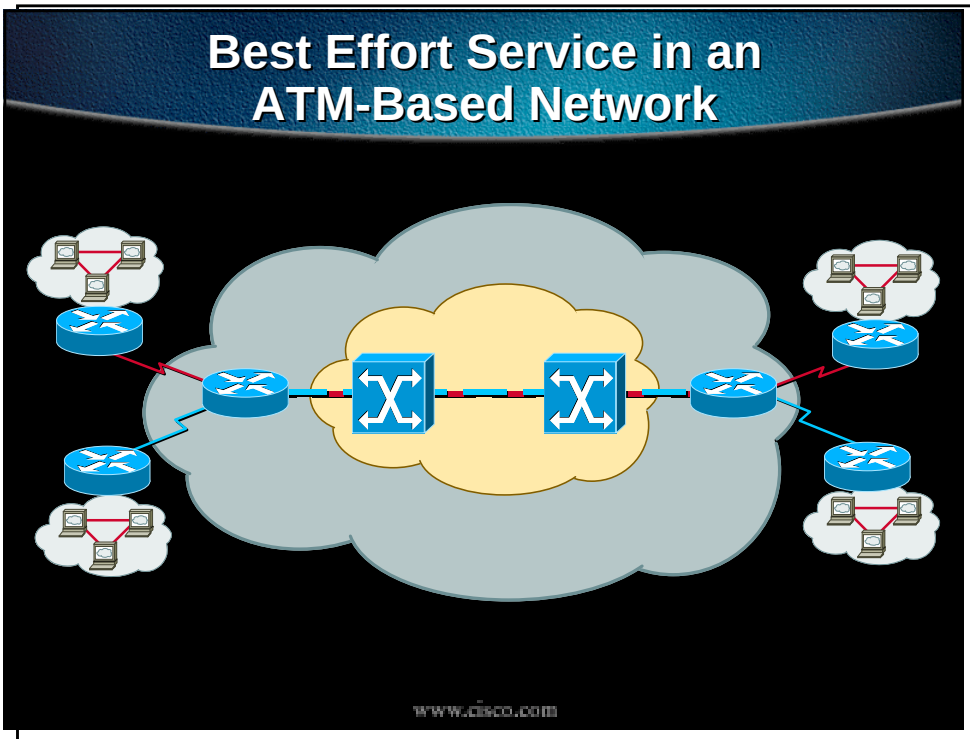

Assured Service in an ATM-Based Network

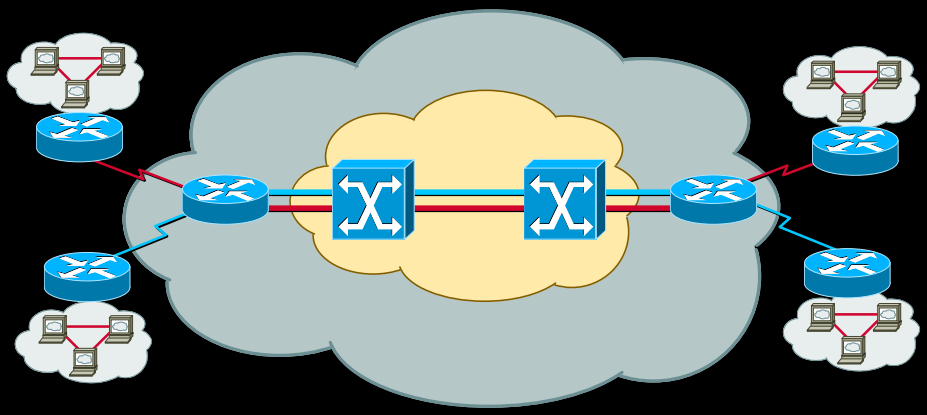

www.cisco.com

So, for TCP

- Traffic can be contained to a rate in a manner consistent with good quality of service
- Traffic can be managed well with a little foresight and planning

www.cisco.com

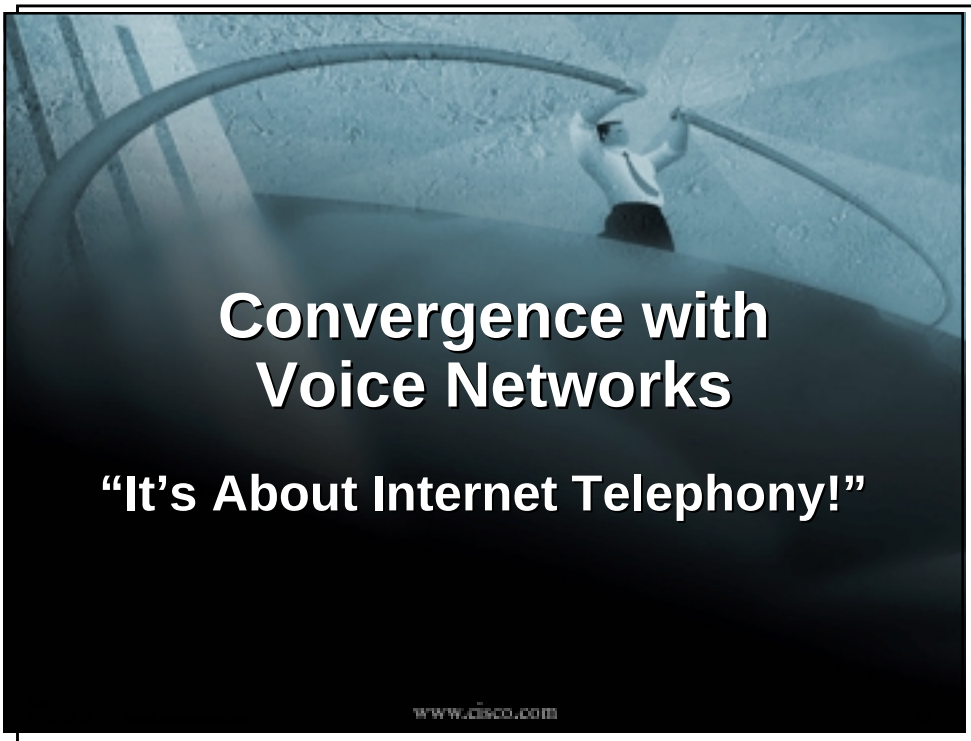

Again, the Premise:

- TCP-based applications, voice, and video can be managed well with a little planning

www.cisco.com

Changing Corporate Network Application Predominance

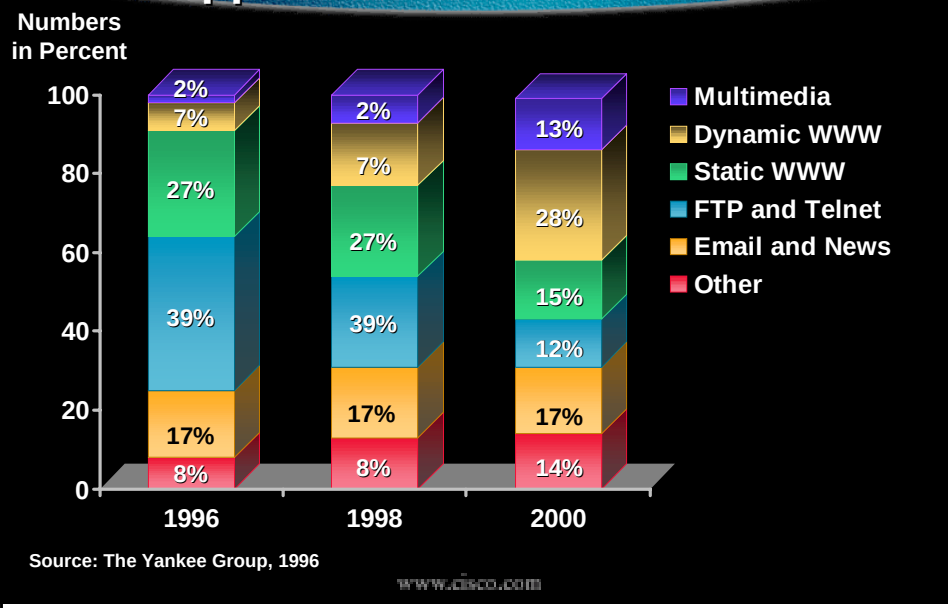

Source: The Yankee Group, 1996

www.cisco.com

Growth of IP Traffic

- Email
- Information search/access
- Subscription services/"Push"
- Conferencing/ multimedia
- Video/imaging

"From 2000 on, 80% of Service Provider Profits Will Be Derived from IP-Based Services."
Source: CIMI Corp.

Source: Multiple IXC Projections

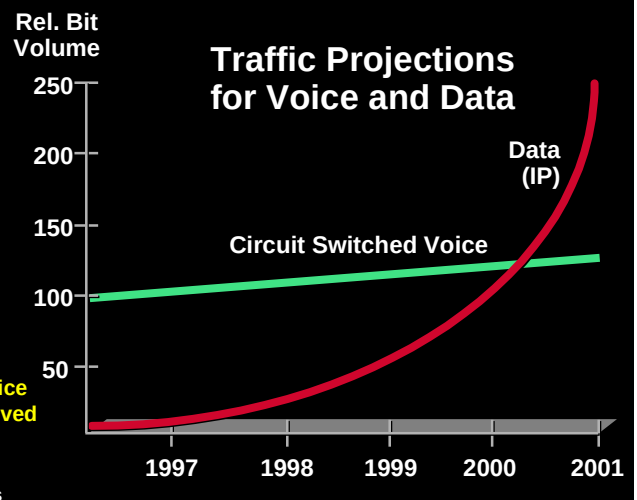

www.cisco.com

High End IP Transport Alternatives

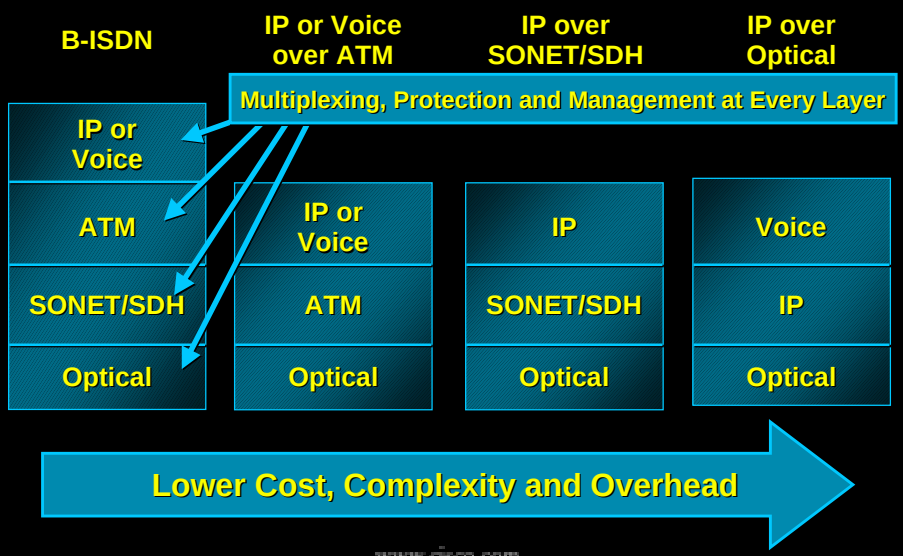

www.cisco.com

H.323 Voice/Video

- Voice
  - Constant bit rate when sending
  - Relatively small messages (44–170 bytes)
- Video
  - Generally high variable bit rate
  - Controlled by codec efficiency on picture
  - Message size is generally the MTU

www.cisco.com

Video: Traffic Pattern

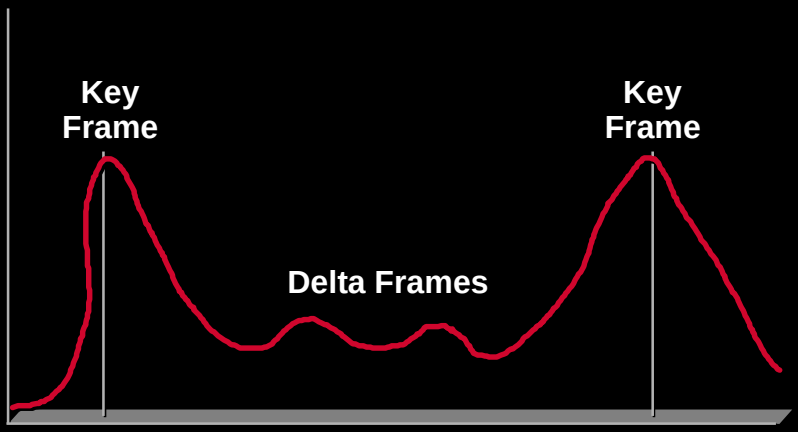

www.cisco.com

Video: Effect of Delay

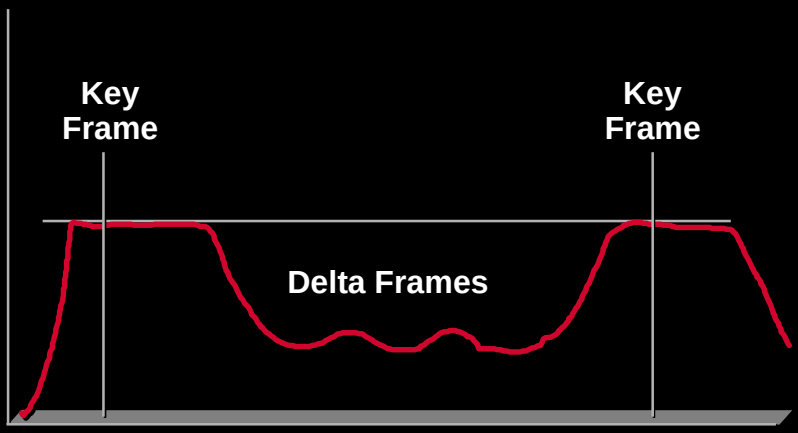

www.cisco.com

Video: Playback Point

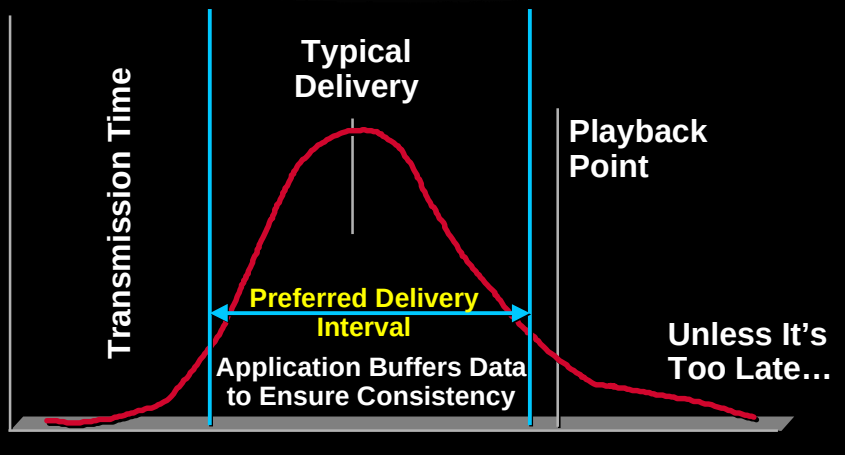

www.cisco.com

Synchronization of Voice and Video

- McGurk effect: voice can sound garbled to human ear when not synchronized with video
- Therefore, we have to synchronize these

www.cisco.com

QoS Definition for Voice

- Low loss rate
- Low absolute delay in two-way situations
  - Broadcast voice doesn't have this problem…
- Low variation in delay

www.cisco.com

Key Issue for Voice QoS

- Silent periods must not be randomly inserted or removed so as to make other sounds unintelligible
- End-to-end delay must be comprehended by human listener

www.cisco.com

QoS Definition for Video

- Low loss rate
- Low absolute delay in two-way situations
- Low variation in delay

www.cisco.com

Key Issues for Video

- All packets that comprise a video frame must arrive during the same frame interval

  OK if it's the last millisecond of that interval…

- Audio and video must be synchronized when shown to user

www.cisco.com

How Can We Make Internet Voice Act Predictably?

- Predictable amount of traffic in the network
- Predictable treatment of packets in routers and switches
- Planning to support these aspects results in a predictable network

Voice/Video Traffic Management Issues

- The fundamental problems with voice/video traffic are:
  - It doesn't slow down in response to delay or loss
  - It requires minimal variation in delay

Predictable Amount of Traffic in the Network

- The implication is that we have to control used capacity
  - Capacity that individual calls consume
    - "If you experience poor quality, use a more compact encoding or a lower frame rate"
  - Capacity that total call volume can consume
    - "If there isn't capacity, refuse new calls"

www.cisco.com

Predictable Treatment of Packets in Routers and Switches

- We have to place voice in queues that give it high priority
  - Maintain tight delay budgets
  - Application of class-based WFQ

www.cisco.com

Planning for a Predictable Network

- Enable CB—WFQ on all relevant links
  - Configure voice queue with more bandwidth than traffic will need, or
  - For low bandwidth, priority queue [12.0(6)T]
- Low speed links should use
  - Link fragmentation or FRF.12
  - RTP compression for voice
- Enable RSVP call negotiation
  - “Refuse excess calls”

www.cisco.com

FRF.12, and Link Fragmentation and Interleaving

- Premise:
  - Reducing voice packet size reduces session requirements on network
  - So compress out IP, UDP, and RTP headers as much as possible
- Limits jitter on lower bandwidth links

www.cisco.com

Resource Reservation

- Current deployment
- Current extensions
- Extensions being developed

www.cisco.com

Current Deployment

- RSVP version 1
  - Call control for individual sessions
  - Deployed
    - Cisco 11.2
    - Microsoft Windows ’98 (service pack)
    - Microsoft Windows NT 2000
- Appropriate to edge networks

www.cisco.com

Current Extensions

- Policy management via COPS
- LAN management via subnet bandwidth manager

www.cisco.com

Policy Management via COPS

- Local or central policy server can authorize decisions
- Local policy:
  - Simple policies
- Central policy server:
  - Certificates,
  - Complex policies

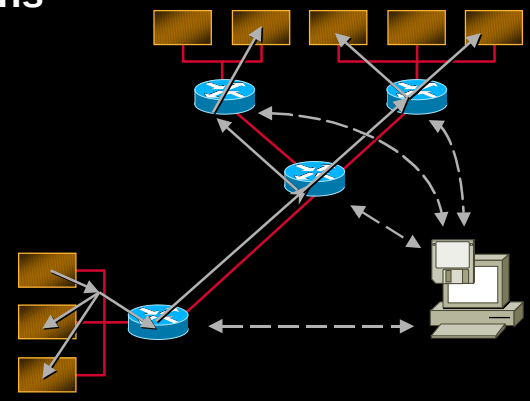

www.cisco.com

LAN Management via Subnet Bandwidth Manager

- Subnet bandwidth manager is RSVP in a switch
- Controls aggregate reservations on a LAN

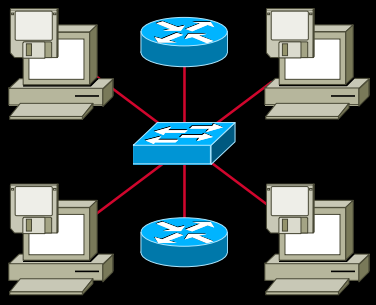

www.cisco.com

Extensions Being Developed

- Rapid deployment of calls
- Aggregate classification in edge networks
- Aggregate classification and admission in service provider networks

www.cisco.com

Rapid Deployment of Calls

- Problem: need acknowledged reservation installation
- Solution: acknowledge it…

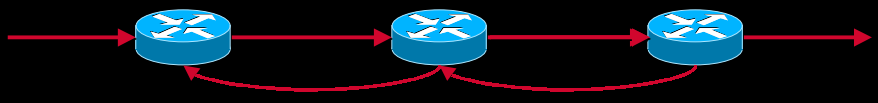

www.cisco.com

Aggregate Classification in Edge Networks

- Use differentiated services code points to identify traffic
  - Rather than specific flows

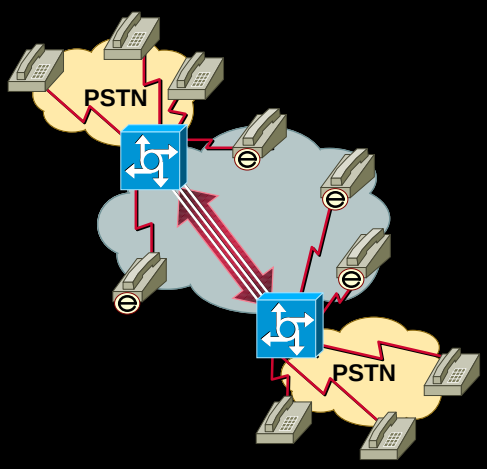

www.cisco.com

Aggregate Classification in Edge Networks

- Reservation requested by host in the usual way (RFC 2205)
- Flow classification and policing at first hop router
- Flow admission along end-to-end path
- Aggregate classification and policing at subsequent routers

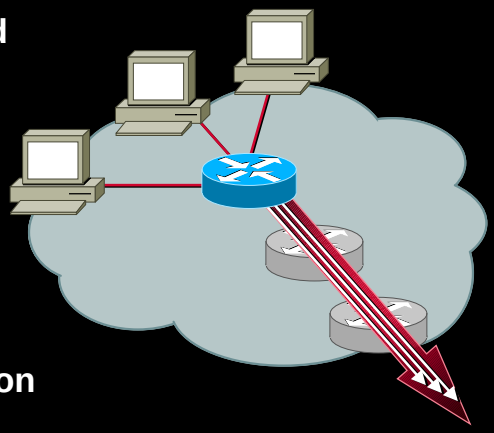

www.cisco.com

Aggregate Classification and Admission Across Service Provider Networks

- Voice/video calls
  - Placed across aggregation domain boundary
- Aggregate reservations
  - Placed from ingress to egress for DSCP used
  - Use expedited forwarding service
  - Limited rate of change

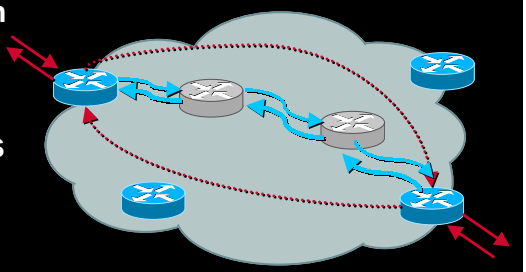

- Why?
  - Otherwise, you don't know that bandwidth exists on a path

www.cisco.com

Solving Voice/Video Issues Using the Expedited Forwarding Service

- Rate control
  - Application at source
  - Reservation in network
- Jitter control
  - WFQ's priority queue (low speed)
  - Statistically empty queue (CB-WFQ)

www.cisco.com

The Implications for Voice and Video

- We can control call volume
  - And therefore traffic volume
- We can scalably prioritize traffic in the system
  - And therefore deliver on latency issues
- So, voice and video can be managed well with a little planning

www.cisco.com

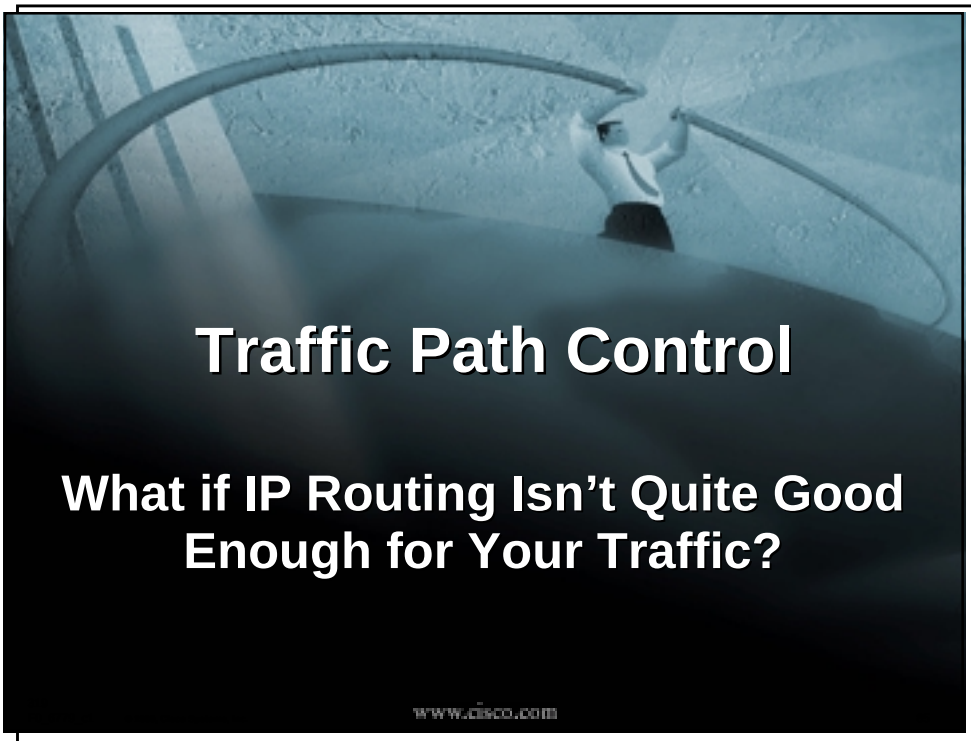

Traffic Path Control

What if IP Routing Isn't Quite Good Enough for Your Traffic?

www.cisco.com

Traffic Engineering

- Historical approaches
  - Load sharing
  - Routing metrics
- A new one
  - Label switching

www.cisco.com

Load Sharing

- Multipath routing
  - Equal and unequal cost
- Multilink PPP

www.cisco.com

Routing

- Administrative metrics
  - Designed to move traffic to statistically low volume links
- Load sensitive metrics
  - Designed to move data away from congested links
  - Tendency towards oscillation

www.cisco.com

Utility of These:

- While they basically work, they are
  - Not deterministic, and
  - Tend to be hard to predict

Multiprotocol Label Switching

- MPLS traffic engineering
  - VPNs and general engineering
- MPLS routing for resource reservation
  - In the direction of QoS routing

Principles of Label Switching

- Labeled paths:
  - Multiple enumerated point-to-point relationships between pairs of routers
  - Sets of pair-wise relationships create a labeled tunnel
- Conceptually similar to ATM VCs or Frame Relay DLCs, but
  - Interface independent
  - Used to model network layer constructs
  - Variable length packets

www.cisco.com

Network Layer Constructs…

- Types of traffic streams
  - Destination routes
  - Source-destination routes
  - AS pairs
  - BGP community pairs
- Tunnels can create
  - Any routing that meets engineering needs

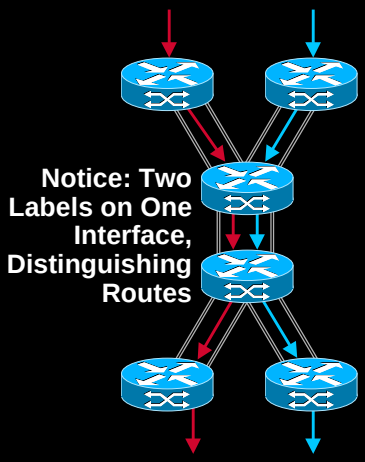

www.cisco.com

Using Labeled Tunnels to Create Virtual Private Networks

- Imagine edge network with private address space
- Stretch labeled tunnels across the network
- Now, do it again
- Disjoint networks
  - Same address space
  - Separate routing

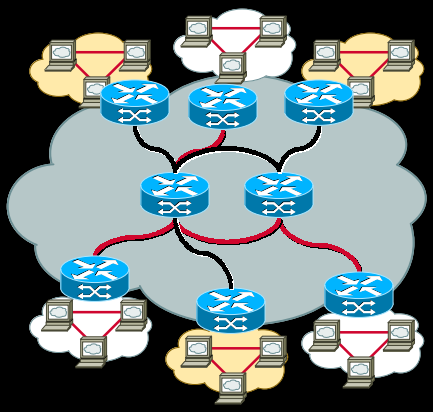

www.cisco.com

MPLS Traffic Engineering

- Same technology can drag specific routes around
  - Several less-used paths vs. a few denser paths…
- Initially seen as off-line engineering
- Can use either LDP or RSVP to install routes

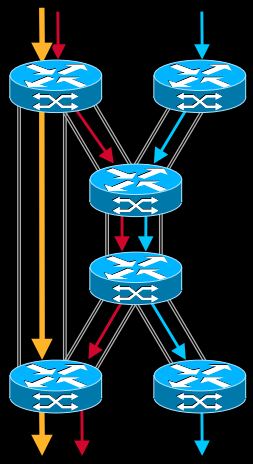

www.cisco.com

CoS in MPLS Networks

- Class of service
  - Roughly similar to diff-serv code point
  - Eight values, not sixty-four
- Implements similar drop/delay management within labeled tunnels
- Therefore, MPLS networks have fundamental TCP QoS support

www.cisco.com

The Obvious Hole…

- Wouldn't it be nice if engineered labeled tunnels could:
  - Have specific bandwidths guaranteed?
  - Recover from network events quickly and automatically using reasonable if not optimal routes?

www.cisco.com

MPLS Routing for Resource Reservation

- Use OSPF/IS-IS to distribute bandwidth availability information
- Edge router does SPF calculation when needed
- RSVP used to install labeled tunnel while checking for race events
- CoS field used to identify traffic for queued rate support

www.cisco.com

Automated Reinstallation of Labeled Tunnels

- RSVP tears down affected tunnels
- Edge devices recalculate routes
- RSVP used to re-install tunnels
- Bandwidth checks result in retry

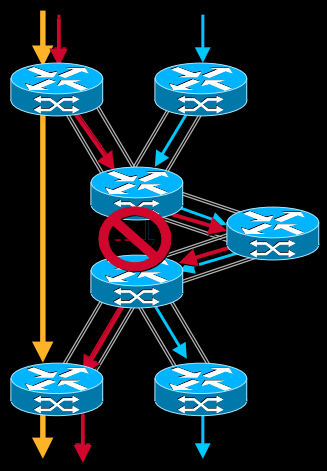

www.cisco.com

Q.E.D. MPLS

- Traffic engineering for network layer traffic can be managed well with a little planning

www.cisco.com

Plan Your Network for Predictability

- Network engineering
- Assured forwarding service
  - TCP
- Expedited forwarding service
  - Voice, implies some form of admission

www.cisco.com

Network Engineering

- Capacity engineering
  - Engineered IP routes?
- May involve traffic engineering
  - Labeled tunnels?

www.cisco.com

Assured Forwarding Service

- Designed for TCP
  - Classes control rates for SLAs
  - Drop controls trace effects back to sources
- Implement using
  - Committed access rate,
  - Weighted Random Early Detection,
  - Class-based weighted fair queuing

www.cisco.com

Expedited Forwarding Service

- Appropriate to voice/video
- Requires
  - Under-subscribed traffic classes
  - Reservation of bandwidth
  - Policing

www.cisco.com

Guiding Principles for Predictability

- One must have at most a predictable amount of traffic in the network
- One must have predictable traffic delay in each network element
- Given these, end-to-end delay of a host to host message is predictable

www.cisco.com

In Your Network…

- TCP-based applications, voice, and video—and your bandwidth—can be managed well with a little planning

www.cisco.com